Network Working Group
Request for Comments: 1183
Updates: RFCs 1034, 1035

C. Everhart
Transarc
L. Mamakos
University of Maryland
R. Ullmann
Prime Computer
P. Mockapetris, Editor
ISI
October 1990

New DNS RR Definitions

Status of this Memo

This memo defines five new DNS types for experimental purposes. This RFC describes an Experimental Protocol for the Internet community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Table of Contents

Introduction..... 1
1. AFS Data Base location..... 2
2. Responsible Person..... 3
2.1. Identification of the guilty party..... 3
2.2. The Responsible Person RR..... 4
3. X.25 and ISDN addresses, Route Binding..... 6
3.1. The X25 RR..... 6
3.2. The ISDN RR..... 7
3.3. The Route Through RR..... 8
REFERENCES and BIBLIOGRAPHY..... 9
Security Considerations..... 10
Authors' Addresses..... 11

Introduction

This RFC defines the format of new Resource Records (RRs) for the Domain Name System (DNS), and reserves corresponding DNS type mnemonics and numerical codes. The definitions are in three independent sections: (1) location of AFS database servers, (2) location of responsible persons, and (3) representation of X.25 and ISDN addresses and route binding. All are experimental.

This RFC assumes that the reader is familiar with the DNS [3,4]. The data shown is for pedagogical use and does not necessarily reflect the real Internet.

1. AFS Data Base location

This section defines an extension of the DNS to locate servers both for AFS (AFS is a registered trademark of Transarc Corporation) and for the Open Software Foundation's (OSF) Distributed Computing Environment (DCE) authenticated naming system using HP/Apollo's NCA, both to be components of the OSF DCE. The discussion assumes that the reader is familiar with AFS [5] and NCA [6].

The AFS (originally the Andrew File System) system uses the DNS to map from a domain name to the name of an AFS cell database server. The DCE Naming service uses the DNS for a similar function: mapping from the domain name of a cell to authenticated name servers for that cell. The method uses a new RR type with mnemonic AFSDB and type code of 18 (decimal).

AFSDB has the following format:

<owner> <ttl> <class> AFSDB <subtype> <hostname>

Both RDATA fields are required in all AFSDB RRs. The <subtype> field is a 16 bit integer. The <hostname> field is a domain name of a host that has a server for the cell named by the owner name of the RR.

The format of the AFSDB RR is class insensitive. AFSDB records cause type A additional section processing for <hostname>. This, in fact, is the rationale for using a new type code, rather than trying to build the same functionality with TXT RRs.

Note that the format of AFSDB in a master file is identical to MX. For purposes of the DNS itself, the subtype is merely an integer. The present subtype semantics are discussed below, but changes are possible and will be announced in subsequent RFCs.

In the case of subtype 1, the host has an AFS version 3.0 Volume Location Server for the named AFS cell. In the case of subtype 2, the host has an authenticated name server holding the cell-root directory node for the named DCE/NCA cell.

The use of subtypes is motivated by two considerations. First, the space of DNS RR types is limited. Second, the services provided are sufficiently distinct that it would continue to be confusing for a client to attempt to connect to a cell's servers using the protocol for one service, if the cell offered only the other service.

As an example of the use of this RR, suppose that the Toaster Corporation has deployed AFS 3.0 but not (yet) the OSF's DCE. Their cell, named toaster.com, has three "AFS 3.0 cell database server"

machines: bigbird.toaster.com, ernie.toaster.com, and henson.toaster.com. These three machines would be listed in three AFSDB RRs. These might appear in a master file as:

```
toaster.com.    AFSDB   1 bigbird.toaster.com.
toaster.com.    AFSDB   1 ernie.toaster.com.
toaster.com.    AFSDB   1 henson.toaster.com.
```

As another example use of this RR, suppose that Femto College (domain name femto.edu) has deployed DCE, and that their DCE cell root directory is served by processes running on green.femto.edu and turquoise.femto.edu. Furthermore, their DCE file servers also run AFS 3.0-compatible volume location servers, on the hosts turquoise.femto.edu and orange.femto.edu. These machines would be listed in four AFSDB RRs, which might appear in a master file as:

```
femto.edu.    AFSDB   2 green.femto.edu.
femto.edu.    AFSDB   2 turquoise.femto.edu.
femto.edu.    AFSDB   1 turquoise.femto.edu.
femto.edu.    AFSDB   1 orange.femto.edu.
```

2. Responsible Person

The purpose of this section is to provide a standard method for associating responsible person identification to any name in the DNS.

The domain name system functions as a distributed database which contains many different form of information. For a particular name or host, you can discover it's Internet address, mail forwarding information, hardware type and operating system among others.

A key aspect of the DNS is that the tree-structured namespace can be divided into pieces, called zones, for purposes of distributing control and responsibility. The responsible person for zone database purposes is named in the SOA RR for that zone. This section describes an extension which allows different responsible persons to be specified for different names in a zone.

2.1. Identification of the guilty party

Often it is desirable to be able to identify the responsible entity for a particular host. When that host is down or malfunctioning, it is difficult to contact those parties which might resolve or repair the host. Mail sent to POSTMASTER may not reach the person in a timely fashion. If the host is one of a multitude of workstations, there may be no responsible person which can be contacted on that host.

The POSTMASTER mailbox on that host continues to be a good contact point for mail problems, and the zone contact in the SOA record for database problem, but the RP record allows us to associate a mailbox to entities that don't receive mail or are not directly connected (namespace-wise) to the problem (e.g., GATEWAY.ISI.EDU might want to point at HOTLINE@BBN.COM, and GATEWAY doesn't get mail, nor does the ISI zone administrator have a clue about fixing gateways).

2.2. The Responsible Person RR

The method uses a new RR type with mnemonic RP and type code of 17 (decimal).

RP has the following format:

<owner> <ttl> <class> RP <mbox-dname> <txt-dname>

Both RDATA fields are required in all RP RRs.

The first field, <mbox-dname>, is a domain name that specifies the mailbox for the responsible person. Its format in master files uses the DNS convention for mailbox encoding, identical to that used for the RNAME mailbox field in the SOA RR. The root domain name (just ".") may be specified for <mbox-dname> to indicate that no mailbox is available.

The second field, <txt-dname>, is a domain name for which TXT RR's exist. A subsequent query can be performed to retrieve the associated TXT resource records at <txt-dname>. This provides a level of indirection so that the entity can be referred to from multiple places in the DNS. The root domain name (just ".") may be specified for <txt-dname> to indicate that the TXT_DNAME is absent, and no associated TXT RR exists.

The format of the RP RR is class insensitive. RP records cause no additional section processing. (TXT additional section processing for <txt-dname> is allowed as an option, but only if it is disabled for the root, i.e., ".").

The Responsible Person RR can be associated with any node in the Domain Name System hierarchy, not just at the leaves of the tree.

The TXT RR associated with the TXT_DNAME contain free format text suitable for humans. Refer to [4] for more details on the TXT RR.

Multiple RP records at a single name may be present in the database. They should have identical TTLs.

EXAMPLES

Some examples of how the RP record might be used.

```
sayshell.umd.edu. A     128.8.1.14
                  MX    10 sayshell.umd.edu.
                  HINFO NeXT UNIX
                  WKS   128.8.1.14 tcp ftp telnet smtp
                  RP    louie.trantor.umd.edu.  LAM1.people.umd.edu.

LAM1.people.umd.edu. TXT (
      "Louis A. Mamakos, (301) 454-2946, don't call me at home!" )
```

In this example, the responsible person's mailbox for the host SAYSHELL.UMD.EDU is louie@trantor.umd.edu. The TXT RR at LAM1.people.umd.edu provides additional information and advice.

```
TERP.UMD.EDU.     A     128.8.10.90
                  MX    10 128.8.10.90
                  HINFO MICROVAX-II UNIX
                  WKS   128.8.10.90 udp domain
                  WKS   128.8.10.90 tcp ftp telnet smtp domain
                  RP    louie.trantor.umd.edu. LAM1.people.umd.edu.
                  RP    root.terp.umd.edu. ops.CS.UMD.EDU.

TRANTOR.UMD.EDU. A     128.8.10.14
                 MX    10 trantor.umd.edu.
                 HINFO MICROVAX-II UNIX
                 WKS   128.8.10.14 udp domain
                 WKS   128.8.10.14 tcp ftp telnet smtp domain
                 RP    louie.trantor.umd.edu. LAM1.people.umd.edu.
                 RP    petry.netwolf.umd.edu. petry.people.UMD.EDU.
                 RP    root.trantor.umd.edu. ops.CS.UMD.EDU.
                 RP    gregh.sunset.umd.edu.  .

LAM1.people.umd.edu.  TXT   "Louis A. Mamakos (301) 454-2946"
petry.people.umd.edu. TXT   "Michael G. Petry (301) 454-2946"
ops.CS.UMD.EDU.       TXT   "CS Operations Staff (301) 454-2943"
```

This set of resource records has two hosts, TRANTOR.UMD.EDU and TERP.UMD.EDU, as well as a number of TXT RRs. Note that TERP.UMD.EDU and TRANTOR.UMD.EDU both reference the same pair of TXT resource records, although the mail box names (root.terp.umd.edu and root.trantor.umd.edu) differ.

Here, we obviously care much more if the machine flakes out, as we've specified four persons which might want to be notified of problems or other events involving TRANTOR.UMD.EDU. In this example, the last RP

RR for TRANTOR.UMD.EDU specifies a mailbox (gregh.sunset.umd.edu), but no associated TXT RR.

3. X.25 and ISDN addresses, Route Binding

This section describes an experimental representation of X.25 and ISDN addresses in the DNS, as well as a route binding method, analogous to the MX for mail routing, for very large scale networks.

There are several possible uses, all experimental at this time. First, the RRs provide simple documentation of the correct addresses to use in static configurations of IP/X.25 [11] and SMTP/X.25 [12].

The RRs could also be used automatically by an internet network-layer router, typically IP. The procedure would be to map IP address to domain name, then name to canonical name if needed, then following RT records, and finally attempting an IP/X.25 call to the address found. Alternately, configured domain names could be resolved to identify IP to X.25/ISDN bindings for a static but periodically refreshed routing table.

This provides a function similar to ARP for wide area non-broadcast networks that will scale well to a network with hundreds of millions of hosts.

Also, a standard address binding reference will facilitate other experiments in the use of X.25 and ISDN, especially in serious inter-operability testing. The majority of work in such a test is establishing the n-squared entries in static tables.

Finally, the RRs are intended for use in a proposal [13] by one of the authors for a possible next-generation internet.

3.1. The X25 RR

The X25 RR is defined with mnemonic X25 and type code 19 (decimal).

X25 has the following format:

<owner> <ttl> <class> X25 <PSDN-address>

<PSDN-address> is required in all X25 RRs.

<PSDN-address> identifies the PSDN (Public Switched Data Network) address in the X.121 [10] numbering plan associated with <owner>. Its format in master files is a <character-string> syntactically identical to that used in TXT and HINFO.

The format of X25 is class insensitive. X25 RRs cause no additional section processing.

The <PSDN-address> is a string of decimal digits, beginning with the 4 digit DNIC (Data Network Identification Code), as specified in X.121. National prefixes (such as a 0) MUST NOT be used.

For example:

```
Relay.Prime.COM.  X25      311061700956
```

3.2. The ISDN RR

The ISDN RR is defined with mnemonic ISDN and type code 20 (decimal).

An ISDN (Integrated Service Digital Network) number is simply a telephone number. The intent of the members of the CCITT is to upgrade all telephone and data network service to a common service.

The numbering plan (E.163/E.164) is the same as the familiar international plan for POTS (an un-official acronym, meaning Plain Old Telephone Service). In E.166, CCITT says "An E.163/E.164 telephony subscriber may become an ISDN subscriber without a number change."

ISDN has the following format:

```
<owner> <ttl> <class> ISDN <ISDN-address> <sa>
```

The <ISDN-address> field is required; <sa> is optional.

<ISDN-address> identifies the ISDN number of <owner> and DDI (Direct Dial In) if any, as defined by E.164 [8] and E.163 [7], the ISDN and PSTN (Public Switched Telephone Network) numbering plan. E.163 defines the country codes, and E.164 the form of the addresses. Its format in master files is a <character-string> syntactically identical to that used in TXT and HINFO.

<sa> specifies the subaddress (SA). The format of <sa> in master files is a <character-string> syntactically identical to that used in TXT and HINFO.

The format of ISDN is class insensitive. ISDN RRs cause no additional section processing.

The <ISDN-address> is a string of characters, normally decimal digits, beginning with the E.163 country code and ending with the DDI if any. Note that ISDN, in Q.931, permits any IA5 character in the

general case.

The <sa> is a string of hexadecimal digits. For digits 0-9, the concrete encoding in the Q.931 call setup information element is identical to BCD.

For example:

```
Relay.Prime.COM.    IN   ISDN      150862028003217
sh.Prime.COM.       IN   ISDN      150862028003217 004
```

(Note: "1" is the country code for the North American Integrated Numbering Area, i.e., the system of "area codes" familiar to people in those countries.)

The RR data is the ASCII representation of the digits. It is encoded as one or two <character-string>s, i.e., count followed by characters.

CCITT recommendation E.166 [9] defines prefix escape codes for the representation of ISDN (E.163/E.164) addresses in X.121, and PSDN (X.121) addresses in E.164. It specifies that the exact codes are a "national matter", i.e., different on different networks. A host connected to the ISDN may be able to use both the X25 and ISDN addresses, with the local prefix added.

3.3. The Route Through RR

The Route Through RR is defined with mnemonic RT and type code 21 (decimal).

The RT resource record provides a route-through binding for hosts that do not have their own direct wide area network addresses. It is used in much the same way as the MX RR.

RT has the following format:

<owner> <ttl> <class> RT <preference> <intermediate-host>

Both RDATA fields are required in all RT RRs.

The first field, <preference>, is a 16 bit integer, representing the preference of the route. Smaller numbers indicate more preferred routes.

<intermediate-host> is the domain name of a host which will serve as an intermediate in reaching the host specified by <owner>. The DNS RRs associated with <intermediate-host> are expected to include at

least one A, X25, or ISDN record.

The format of the RT RR is class insensitive. RT records cause type X25, ISDN, and A additional section processing for <intermediate-host>.

For example,

```
sh.prime.com.       IN   RT   2    Relay.Prime.COM.
                    IN   RT   10   NET.Prime.COM.
*.prime.com.        IN   RT   90   Relay.Prime.COM.
```

When a host is looking up DNS records to attempt to route a datagram, it first looks for RT records for the destination host, which point to hosts with address records (A, X25, ISDN) compatible with the wide area networks available to the host. If it is itself in the set of RT records, it discards any RTs with preferences higher or equal to its own. If there are no (remaining) RTs, it can then use address records of the destination itself.

Wild-card RTs are used exactly as are wild-card MXs. RT's do not "chain"; that is, it is not valid to use the RT RRs found for a host referred to by an RT.

The concrete encoding is identical to the MX RR.

REFERENCES and BIBLIOGRAPHY

[1] Stahl, M., "Domain Administrators Guide", RFC 1032, Network Information Center, SRI International, November 1987.

[2] Lottor, M., "Domain Administrators Operations Guide", RFC 1033, Network Information Center, SRI International, November, 1987.

[3] Mockapetris, P., "Domain Names - Concepts and Facilities", RFC 1034, USC/Information Sciences Institute, November 1987.

[4] Mockapetris, P., "Domain Names - Implementation and Specification", RFC 1035, USC/Information Sciences Institute, November 1987.

[5] Spector A., and M. Kazar, "Uniting File Systems", UNIX Review, 7(3), pp. 61-69, March 1989.

[6] Zahn, et al., "Network Computing Architecture", Prentice-Hall, 1989.

[7] International Telegraph and Telephone Consultative Committee,

"Numbering Plan for the International Telephone Service", CCITT Recommendations E.163., IXth Plenary Assembly, Melbourne, 1988, Fascicle II.2 ("Blue Book").

[8] International Telegraph and Telephone Consultative Committee, "Numbering Plan for the ISDN Era", CCITT Recommendations E.164., IXth Plenary Assembly, Melbourne, 1988, Fascicle II.2 ("Blue Book").

[9] International Telegraph and Telephone Consultative Committee. "Numbering Plan Interworking in the ISDN Era", CCITT Recommendations E.166., IXth Plenary Assembly, Melbourne, 1988, Fascicle II.2 ("Blue Book").

[10] International Telegraph and Telephone Consultative Committee, "International Numbering Plan for the Public Data Networks", CCITT Recommendations X.121., IXth Plenary Assembly, Melbourne, 1988, Fascicle VIII.3 ("Blue Book"); provisional, Geneva, 1978; amended, Geneva, 1980, Malaga-Torremolinos, 1984 and Melborne, 1988.

[11] Korb, J., "Standard for the Transmission of IP datagrams Over Public Data Networks", RFC 877, Purdue University, September 1983.

[12] Ullmann, R., "SMTP on X.25", RFC 1090, Prime Computer, February 1989.

[13] Ullmann, R., "TP/IX: The Next Internet", Prime Computer (unpublished), July 1990.

[14] Mockapetris, P., "DNS Encoding of Network Names and Other Types", RFC 1101, USC/Information Sciences Institute, April 1989.

Security Considerations

Security issues are not addressed in this memo.

Authors' Addresses

Craig F. Everhart
Transarc Corporation
The Gulf Tower
707 Grant Street
Pittsburgh, PA 15219

Phone: +1 412 338 4467

EMail: Craig_Everhart@transarc.com

Louis A. Mamakos
Network Infrastructure Group
Computer Science Center
University of Maryland
College Park, MD 20742-2411

Phone: +1-301-405-7836

Email: louie@Sayshell.UMD.EDU

Robert Ullmann 10-30
Prime Computer, Inc.
500 Old Connecticut Path
Framingham, MA 01701

Phone: +1 508 620 2800 ext 1736

Email: Ariel@Relay.Prime.COM

Paul Mockapetris
USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292

Phone: 213-822-1511

EMail: pvm@isi.edu